arXiv:1606.06271v2 [cs.GT] 4 Aug 2016

Dynamic Programming for One-Sided Partially Observable Pursuit-Evasion Games

Karel Horák, Branislav Bošanský
{horak,bosansky}@agents.fel.cvut.cz

Department of Computer Science, Faculty of Electrical Engineering,
Czech Technical University in Prague

Abstract. Pursuit-evasion scenarios appear widely in robotics, security domains, and many other real-world situations. We focus on two-player pursuit-evasion games with concurrent moves, infinite horizon, and discounted rewards. We assume that the players have a partial observability, however, the evader is given an advantage of knowing the current position of the units of the pursuer. This setting is particularly interesting for security domains where a robust strategy, designed to maximize the utility in the worst-case scenario, is often desirable. We provide, to the best of our knowledge, the first algorithm that provably converges to the value of a partially observable pursuit-evasion game with infinite horizon. Our algorithm extends well-known value iteration algorithm by exploiting that (1) the value functions of our game depend only on position of the pursuer and the belief he has about the current position of the evader, and (2) that these functions are piecewise linear and convex in the belief space.

1 Introduction

Pursuit-evasion games appear in many scenarios in robotics and security domains [13,2], where a team of centrally controlled pursuing units (the *pursuer*) aims to locate and capture the *evader*, while the evader aims for the opposite. We study this class of games and assume their discrete-time variant played on a finite graph. We further assume that all units of both players move simultaneously, that the horizon of the game is infinite, the rewards are discounted over time with discount factor $\gamma \in [0, 1)$, and that the players have only a partial information about the state of the world. Formally, such a game belongs to zero-sum partially observable stochastic games (POSGs).

We are interested in finding robust strategies of the pursuer against the worst-case evader. Specifically, we assume that the evader knows the positions of the pursuing units and her only uncertainty is the strategy of the pursuer and the move that will be performed in the current time step. Although in reality such perfectly informed adversary is rarely met, it is typical that the pursuer does not know what information is being revealed to the evader. Hence, in order to derive robust strategies (i.e. those that maximize pursuer's reward against *any* type of the evader), it is natural to consider this kind of perfectly informed adversary.

We design the first algorithm that provably converges to the value of such one-sided partially observable pursuit-evasion games. Moreover as the value converges, the strategies of the players converge to their optimal strategies as well. This is in contrast to the existing approaches in robotics and security, where heuristic solutions without any optimality guarantees are used [13,2].

Our algorithm extends the well-known value iteration algorithm that is known to work for concurrent-moves stochastic games [9] as well as for partially observable models from decision theory – Partially Observable Markov Decision Processes (POMDPs) [10,7,8,11]. We adopt the methodology for POMDPs and show that one-sided pursuit-evasion games also allow us to define compactly represented value functions and thus to design a dynamic programming algorithm that iteratively improves values of game states over time and converges to the value of the game. Specifically we show that the value functions (1) depend only on the position of the units of the pursuer and the belief he has about the possible position of the evader, but do *not* depend on the history of moves, (2) these functions are piecewise linear and convex and thus can be, similarly to POMDPs, represented as a set of so called α-vectors (Section 2.1), and (3) we can design a dynamic-programming operator that provably converges to optimal value of the game when applied on these value functions (Section 3).

We believe that our algorithm (accompanied with related theoretical results) is of comparable significance in this class of games as the full-backed value iteration algorithm in the class of POMDPs and it is a necessary first step towards designing practical scalable algorithms.

Due to the space constraints, most of the technical proofs can be found in the Appendix.

1.1 Related Work

A similar model with one-sided partial observability where one of the players has a perfect information was presented by McEneaney [6]. The author assumed that the player with perfect information knows the action the opponent plays at the current stage. Due to the turn-based character of such game, the author considers only pure strategies.

The difference between our concurrent setup and the turn-based one is highlighted in Fig. 1. In the turn-based setting, the evader always observes the action played by the pursuer before making her action. This allows the evader to always move to a vertex that is unoccupied by the pursuer, hence preventing the pursuer from ever capturing him. On the other hand, in the concurrent setting, the only thing the evader can get to know is that the pursuer chooses every vertex with equal probability. She cannot anticipate the move of the pursuer and hence gets captured with probability 1/3 in the first round of the game.

Our setting with concurrent moves better corresponds to the real-world situations that occur in real time. As the evader does not know the action taken by the pursuer in the current stage, players may need to use randomized strategies. However, allowing randomized strategies provides challenges in the design of the dynamic programming operator that we address in this paper.

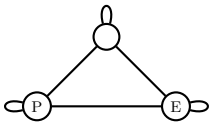

Fig. 1. A game with one unit of the pursuer (node marked with P) and the evader (node marked with E). Pure strategies are insufficient if players act simultaneously.

Another model that uses one-sided partial observability was considered by Chatterjee et al. [1], however with reachability and safety objectives (a player either wants to reach a set of target states or she wants to keep the system in a set of safe states) that do not translate to objectives with discounted rewards.

An algorithm for solving a broader class of POSGs, where all players have imperfect information, was proposed in [4], yet considering the finite horizon only. The information is not shared among players, hence they typically attain different beliefs about the state of the game. The algorithm uses ideas of dynamic programming to incrementally construct a set of pure strategies, longer ones from the shorter ones. The set of such strategies is then used to form a normal-form representation of the POSG. The use of iterated elimination of dominated strategies ensures that irrelevant strategies are not considered. Nevertheless, in some games all strategies are relevant and this elimination does not help. This is the case of the game in Fig. 1 where the uniform play of the pursuer involves randomization over *all* pure strategies considered by the algorithm; the number of such strategies being exponential in the horizon. The one-sided partial observability in our game allows us to avoid such enumeration of pure strategies, define dynamic programming over values of the subgames, and, in this particular case, represent the optimal *infinite* horizon strategy using a single α-vector.

2 Finite-horizon game

We use the notion of finite-horizon POSGs, or *extensive-form games*, to reason about the infinite-horizon pursuit-evasion game with discounted rewards. An extensive-form game (EFG) is a tuple $G = (\mathcal{N}, \mathcal{H}, \mathcal{Z}, \mathcal{T}, u, \mathcal{I})$. $\mathcal{N}$ is the set of players, in our case $\mathcal{N} = \{p, e\}$ where p stands for the pursuer and e for the evader. Set $\mathcal{H}$ denotes a finite set of *histories* of actions taken by all players from the begining of the game. Every history corresponds to a *node* in the game tree; hence we use the terms history and node interchangeably. Each of the histories may be either (1) *terminal* ($h \in \mathcal{Z} \subseteq \mathcal{H}$) where the game ends and player i gets utility $u_i(h)$, (2) controlled by the nature player who selects the successor node according to a fixed probability distribution known to all players, or (3) one of the players from $\mathcal{N}$ may be to act. We consider a zero-sum scenario where $u_p(h) = -u_e(h)$. To simplify the notation we use $u(h)$ to denote pursuer's reward. An ordered list of transitions of player i from root to node h is referred to as a player i's *sequence*. The allowed transitions in the game are modelled using a *transition function* $\mathcal{T}$ that provides a set of successor nodes for each non-terminal history. The imperfect observation of players is modelled via

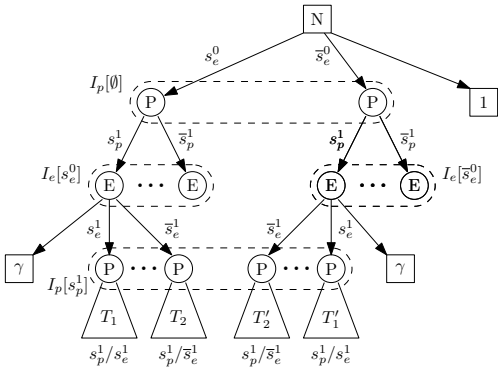

Fig. 2. An EFG representation of a finite-horizon pursuit-evasion game.

information sets $\mathcal{I}_i$ that form a partition over histories h where player $i \in \mathcal{N}$ takes action. We assume perfect recall setting where the players never forget their past actions, i.e. for every $I_i \in \mathcal{I}_i$, all histories $h \in I_i$ have the same player i's sequence. Each information set $I_i \in \mathcal{I}_i$ corresponds to one decision point of player i. A randomized *behavioral strategy* of player i assigns a distribution over actions to each of the information sets in $\mathcal{I}_i$. A behavioral strategy of player i can be represented in the form of a *realization plan* r which assigns probability of playing sequence σ_i to each player i's sequence σ_i. The behavioral strategy at information set $I_i \in \mathcal{I}_i$ reached using a sequence σ_i is then $b(I_i, a) = r(\sigma_i a)/r(\sigma_i)$. A Nash equilibrium (NE) in an EFG is a pair of behavioral strategies, in which each player plays a best response to the strategy of the opponent. The expected utility of the pursuer when NE strategies are played by both players is referred to as the *value of the game*.

We will now use this terminology to construct an EFG for a finite-horizon version of a pursuit-evasion game with N pursuing units played on a graph $\mathcal{G} = \langle \mathcal{V}, \mathcal{E} \rangle$ for t rounds (we term t as the *horizon*). Part of the game tree is shown in Fig. 2. At every round $\tau \leq t$, pursuer's units occupy vertices s_p^τ, where $s_p^\tau = \left\{ s_{p,1}^\tau, \ldots, s_{p,N}^\tau \right\}$ is an N-element multiset of vertices of $\mathcal{G}$, and the evader is located in vertex $s_e^\tau \in \mathcal{V}$. The goal of the pursuer is to achieve a situation where the evader is caught, i.e. $s_e^\tau \in s_p^\tau$. In every round, players have to move their units to vertices adjacent to their current positions ($\mathrm{adj}(v)$ denotes the set of vertices adjacent to v). Position of the evader in round $\tau + 1$ is thus $s_e^{\tau+1} \in \mathrm{adj}(s_e^\tau)$. We overload the operator adj to apply it also on multisets representing positions of pursuer's units, i.e. $s_p^{\tau+1} \in \mathrm{adj}(s_p^\tau)$, where $\mathrm{adj}(s_p^\tau) = \times_{i=1\ldots N} \mathrm{adj}(s_{p,i}^\tau)$.

A horizon-t game $G^t\langle s_p^0, b^0 \rangle$ is parametrized by the initial position of the pursuer $s_p^0 \in \mathcal{V}^N$ and a distribution over evader's initial positions $b^0 \in \Delta(\mathcal{V})$ known to both players (we term b^0 the initial *belief*). The game starts with a chance move selecting the initial position of the evader s_e^0 (based on b^0).

A history $h \in \mathcal{H}$ in a game with horizon t corresponds to a list of positions $s_e^0 s_p^1 s_e^1 \cdots s_p^\tau s_e^\tau$, where $\tau \leq t$. The utility values are assigned to terminal histories as follows: in case the pursuer failed to capture the evader in time, i.e. if $\tau = t$

and $s_e^\tau \notin s_p^\tau$, the pursuer gets utility $u(h) = 0$; if he successfully captured the evader in the time limit t, i.e. if $\tau \leq t$ and $s_e^\tau \in s_p^\tau$, the pursuer gets the reward $u(h) = \gamma^\tau$ for capturing the evader in τ rounds (where $\gamma \in [0, 1)$ stands for the *discount factor*). The transition function $\mathcal{T}$ complies with the graph (i.e., the adjacency function adj), hence $s_p^\tau \in \mathrm{adj}(s_p^{\tau-1})$ and $s_e^\tau \in \mathrm{adj}(s_e^{\tau-1})$ for every $\tau \geq 1$. For notational simplicity we denote the sequence of pursuer's actions $s_p^1 \cdots s_p^\tau$ in h as $h|_p$ and the sequence of evader's actions $s_e^1 \cdots s_e^\tau$ in h as $h|_e$.

The position of the evader is unknown to the pursuer. Hence, in a perfect recall game, there is one pursuer's information set $I_p[\sigma_p]$ for each of his sequences σ_p where $I_p[\sigma_p] = \{h' \mid h' \in \mathcal{H} \setminus \mathcal{Z} \colon h'|_p = \sigma_p\}$.

Evader on the other hand knows the game situation almost perfectly. She knows where the pursuer's units were located *before* the pursuer acted in the current round of the game (recall that we assume that pursuer acts first). The only information missing to the evader is the action being taken by the pursuer in the current round. Hence, for every history $h = s_e^0 s_p^1 s_e^1 \cdots s_p^\tau s_e^\tau$ where the pursuer is to play, there is evader's information set $I_e[h] = \left\{h s_p^{\tau+1} | s_p^{\tau+1} \in \mathrm{adj}(s_p^\tau)\right\}$ containing all possible continuations of the pursuer.

2.1 Shape of the value function

The size of the extensive-form game representation and associated behavioral strategies grows exponentially as the horizon increases. This makes it quickly impossible to apply standard algorithms operating on game trees, especially since we aim to solve *infinite* horizon games.

We aleviate the problem of increasing complexity of the strategy representation by representing them *only* as their values. We show that the value of a strategy is linear in the belief, and we can thus represent it using just $|\mathcal{V}|$ real numbers. Moreover we show that there is a finite set of behavioral strategies that needs to be considered in a game with arbitrary finite horizon, regardless of the initial belief, and thus the value function representing the value of the best strategy at every belief is piecewise linear and convex, which allows us to represent this function in a compact manner.

Definition 1. *A* value function $v^t\langle s_p^0 \rangle : \Delta(\mathcal{V}) \to [0, 1]$ *is a function assigning the value* $v^t\langle s_p^0 \rangle (b^0)$ *of the game* $G^t\langle s_p^0, b^0 \rangle$ *to every initial belief* b^0 *about the position of the evader. By* v^t *we mean a set of value functions* $v^t\langle s_p^0 \rangle$*, one for each initial position* $s_p^0 \in \mathcal{V}^N$ *of the pursuer.*

In the following text we show that a value function $v^t\langle s_p^0 \rangle$ is piecewise linear and convex (PWLC) in the belief for every finite horizon t. For notational simplicity, the term linear will be used to refer to an affine function as well. The proof is structured as follows: (1) first of all we show that the expected utility of every strategy of the pursuer is linear in the belief, next (2) there is a finite set of behavioral strategies $\Sigma^t\langle s_p^0 \rangle$ applicable in games $G^t\langle s_p^0, \cdot \rangle$ such that for every initial belief $b^0 \in \Delta(\mathcal{V})$ the pursuer has about the position of the evader, at least

one strategy in $\Sigma^t\langle s_p^0\rangle$ is a NE solution of the game $G^t\langle s_p^0, b^0\rangle$; and finally (3) we show that the PWLC nature of the value function follows from (1) and (2).

Lemma 1. *Let σ_p be a randomized behavioral strategy of the pursuer in games $G^t\langle s_p^0, b^0\rangle$, where the pursuer starts in vertices s_p^0, parametrized by the initial belief about the position of the evader b^0. The expected utility of playing σ_p against a best responding opponent is linear in b^0.*

Proof. As the strategy σ_p of the pursuer is fixed and the actual evader's position s_e^0 is revealed to the evader, it is not necessary from her perspective to consider the initial belief b^0 to derive the optimal evasive plan. She therefore chooses the optimal evasive plan (minimizing expected pursuer's utility) for every her possible initial position, value of each such plan being a constant. The initial belief b^0 forms a convex combination of values of individual evasive plans, which is a linear function in the belief space. □

Theorem 1. *Let $G^t\langle s_p^0, b^0\rangle$ be a horizon-t game parametrized by the initial belief b^0 where the pursuer starts in a set of vertices s_p^0. There exists a finite set of pursuer's behavioral strategies $\Sigma^t\langle s_p^0\rangle$ such that for every initial belief b^0 about the position of the evader, the set $\Sigma^t\langle s_p^0\rangle$ contains at least one strategy $\sigma_p \in \Sigma^t\langle s_p^0\rangle$ that is in Nash equilibrium of $G^t\langle s_p^0, b^0\rangle$.*

Proof. We use the sequence-form linear program for solving EFGs [5] to reason about the set of strategies $\Sigma^t\langle s_p^0\rangle$ the pursuer has to consider. In this LP, values in every information set of the evader, as well as the value $v(root)$ in the root node of the game tree, are computed in a bottom-up fashion. Every such value $v(I_e)$ of an information set I_e can be seen as a concave piecewise linear function in the space of pursuer's realization plans (a compact representation of his behavioral strategies). The pursuer then seeks for a realization plan that maximizes $v(root)$; the maximizer of which can be found among extreme points of line segments of $v(root)$, i.e. vertices of a polytope bounded by this function [12]. We show that the set of such extreme points does not depend on the initial belief b^0.

There is one information set $I_e[s_e^0]$ of the evader for each of her initial positions s_e^0. The utility of every terminal node in the subgame beneath $I_e[s_e^0]$ is multiplied with chance probability $b(s_e^0)$, which allows us to factor out this probability and obtain the following constraint for the root node:

$$v(root) \leq \sum_{s_e^0 \in s_p^0} b^0(s_e^0) + \sum_{s_e^0 \in \mathcal{V} \setminus s_p^0} b^0(s_e^0) \cdot \hat{v}(I_e[s_e^0]) \tag{1}$$

Value $v(root)$ is a convex combination of concave piecewise linear functions $\hat{v}(I_e[s_e^0])$. As the belief was factored out, these functions, as well as the finite set of their extreme points $P[s_e^0]$, no longer depend on the belief. This convex combination with arbitrary coefficients b^0 cannot have an extreme where none of the functions $\hat{v}(I_e[s_e^0])$ has one. The set of extreme points is therefore a subset of $\bigcup_{s_e^0} P[s_e^0]$ — a finite set that does not depend on the belief. Each of the extreme points in $\bigcup_{s_e^0} P[s_e^0]$ corresponds to one pursuer's realization plan, and thus one his behavioral strategy, which allows us to construct the finite set $\Sigma^t\langle s_p^0\rangle$. □

Theorem 2. *Value function $v^t\langle s_p^0\rangle$ is piecewise linear and convex in the belief space.*

Proof. This result directly follows from Lemma 1 and Theorem 1. There is a finite set of randomized strategies $\Sigma^t\langle s_p^0\rangle$ that has to be considered by the pursuer and value of each such strategy is linear in the belief space. Thus the value function $v^t\langle s_p^0\rangle$ is a pointwise maximum taken over a finite set of linear functions, which is a PWLC function in the belief space. □

Every PWLC function can be represented as a finite set of α-*vectors*. Every such α-vector $\alpha = (\alpha_1, \ldots, \alpha_{|\mathcal{V}|})$ represents one of the affine functions by assigning an expected reward α_i to each of the pure beliefs. We will often be working with the α-vector representation of the value function, hence we overload the notation and consider value functions also as sets of such α-vectors.

Lemma 1 and Theorem 1 imply that every linear segment of the PWLC value function corresponds to one randomized strategy of the pursuer. This is similar to the POMDP case where each α-vector corresponds to one conditional plan. This allows us to use terms α-vector and pursuer's strategy interchangeably.

3 Value iteration

In the previous section, we related the concept of the value functions to the EFG representation of the game and proved that these functions have desirable properties (they are piecewise linear and convex). We leverage their representation to design a dynamic programming approach inspired by value iteration algorithms for either POMDPs [10,7] or perfect information stochastic games [9].

The algorithm inductively constructs a sequence of value functions $\{v^t\}_{t=0}^{\infty}$, starting with values of a horizon-0 game (where the utility of the pursuer depends solely on the fact whether the evader starts in a vertex where one of the pursuer's units is located).

We avoid using the exponentially-sized representation of the underlying EFG by computing value function of a horizon-t game using the solution of the game with horizon $t-1$. First of all we show that there is a well-defined value update formula that expresses values of v^t using value functions v^{t-1} (Theorem 3). We let the players choose their strategies for the first round of the horizon-t game using the maximin principle (we term such strategies *one-step strategies*) and we show that the pursuer can use these strategies to update his belief. Pursuer's one-step strategy π_p is a distribution over possible actions of his units, $\pi_p \in \mathrm{adj}(s_p^0)$, from which he samples his action. The evader acts similarly, however she conditions her decision on her *true* position s_e^0 (not just on the overall belief available to the pursuer); her one-step strategy is thus a mapping $\pi_e : \mathcal{V} \to \Delta(\mathcal{V})$, such that $\pi_e(s_e^0)$ assigns zero probability to vertices not adjacent to s_e^0.

The piecewise linearity and convexity of value functions have implications on the computation of value functions. Firstly it allows us to find optimal one-step strategies by means of linear programming (Section 3.1), furthermore it makes it possible to avoid evaluating the value update formula in every point in

the belief space when constructing new value functions. Instead it is possible to construct new value functions incrementally and focus only on beliefs where the function being constructed has its extreme points of line segments (Section 3.2). We conclude by showing that the dynamic programming operator used in the value iteration algorithm has a unique fixpoint corresponding to value functions of an infinite horizon game which ensures convergence properties of the algorithm (Theorem 4).

Theorem 3. *The value of the game* $G^t\langle s_p^0, b\rangle$ *can be computed from the solutions of horizon-*$(t-1)$ *games, whose values are represented by a set of value functions* v^{t-1}*. It holds that*

$$v^t\langle s_p^0\rangle(b) = \sum_{s_e \in s_p^0} b(s_e) + \gamma \left[\sum_{s_e \in \mathcal{V}\setminus s_p^0} b(s_e)\right] \cdot \max_{\pi_p} \min_{\pi_e} \sum_{s_p^1 \in \mathcal{V}^N} \pi_p(s_p^1) \cdot v^{t-1}\langle s_p^1\rangle(b_{\pi_e}) \quad (2)$$

where the transformed belief b_{π_e} *depends solely on the evader's one-step strategy* π_e *and the parametrization of the game* $G^t\langle s_p^0, b\rangle$*:*

$$b_{\pi_e}(s_e') = \frac{1}{\sum_{s_e \in \mathcal{V}\setminus s_p^0} b(s_e)} \sum_{s_e \in \mathcal{V}\setminus s_p^0} b(s_e) \cdot \pi_e(s_e, s_e') \quad (3)$$

The computation of v^t *by means of Equation* (2) *forms a dynamic programming operator* H*, such that* $v^t = Hv^{t-1}$*.*

Proof. The correctness of the value update formula will be proven by computing the value of the game $G^t\langle s_p^0, b\rangle$ in a bottom-up fashion. We start by considering that the one-step strategies of the players for the first round of the game are fixed, while they play optimally in the rest of the game. This determines pursuer's expected reward at every node in the game tree, which allows us to express his expected utility in the root node of the game tree as an expectation over expected rewards in subsequent nodes (Lemma 2). Due to the fixed behavior in the first round of the game, parts of the game tree are independent on each other — we refer to these subgames as $G[s_p^1]$. This allows us to evaluate the expectation depicted in Lemma 2 by solving these games separately. It turns out that games $G[s_p^1]$ are strategically equivalent to a game $G^{t-1}\langle s_p^1, b_{\pi_e}\rangle$ with shorter horizon, the solution of which is represented by value functions v^{t-1}. The expectation can thus be expressed solely in terms of v^{t-1}. Finally the assumption of fixed one-step strategies for the first round of the game gets relaxed, which yields the desired maximin formula from Equation (2). The derivation of the value update formula relies on several technical lemmas, the proofs of which can be found in the Appendix.

Let $\pi_p \in \Delta(\mathrm{adj}(s_p^0))$ be a fixed pursuer's one-step strategy, and $\pi_e : \mathcal{V} \to \Delta(\mathcal{V})$ be a fixed one-step strategy of the evader. Assume that both players play according to π_p and π_e in the first round of the game, i.e. the pursuer follows π_p in his information set $I_p[\emptyset]$ (i.e. pursuer's information set where he has not acted yet, see Fig. 2) and the evader plays according to $\pi_e(s_e^0)$ in her information set $I_e[s_e^0]$ (where she has received the information that she is located in vertex s_e^0). Once the first round of the game is over, players continue with their best strategies for the situation they are currently in. We denote such optimal strategies where the players are restricted to play π_p and π_e in the first round as σ_p and σ_e.

Definition 2. *Let π_p and π_e be fixed one-step strategies of the players for the first round of the game $G^t\langle s_p^0, b\rangle$ and let σ_p, σ_e be optimal strategies of the players with the restriction to play π_p and π_e in the first round. The expected reward of the pursuer when strategies (σ_p, σ_e) are followed and node h in the game tree is reached is denoted $u(h)$ and termed* expected reward in h.

We follow by expressing the expected utility the pursuer gets when strategies (σ_p, σ_e) are followed by propagating expected rewards from subsequent nodes in the game tree. We use histories of the form $s_e^0 s_p^1 s_e^1$ where the evader started in vertex s_e^0 (based on the move of nature) and then, during the first round of the game, the pursuer moved his units to vertices s_p^1 and the evader moved to s_e^1.

Lemma 2. *The expected reward in the root node equals to:*

$$u(\emptyset) = \sum_{s_e^0 \in s_p^0} b(s_e^0) + \left[\sum_{s_e^0 \notin s_p^0} b(s_e^0)\right] \cdot \sum_{s_p^1} \pi_p(s_p^1)\left(\gamma \sum_{s_e^1 \in s_p^1} b_{\pi_e}(s_e^1) + \right.$$
$$\left. + \left[\sum_{s_e^1 \notin s_p^1} b_{\pi_e}(s_e^1)\right] \sum_{s_e^1 \notin s_p^1} \sum_{s_e^0 \notin s_p^0} \left[\frac{b(s_e^0) \cdot \pi_e(s_e^0, s_e^1)}{\sum_{\tilde{s}_e^1 \notin s_p^1} \sum_{\tilde{s}_e^0 \notin s_p^0} b(\tilde{s}_e^0) \cdot \pi_e(\tilde{s}_e^0, \tilde{s}_e^1)} \cdot u(s_e^0 s_p^1 s_e^1)\right]\right) \quad (4)$$

Lemma 2 expressed the value in the root node based on the expected rewards in histories $s_e^0 s_p^1 s_e^1$ where the pursuer is to move. The pursuer knows only s_p^1, hence these histories are partitioned into his information sets $I_p[s_p^1]$, one for each move s_p^1 of the pursuer in the first round (see Fig. 2). Importantly, for every subgame below $I_p[s_p^1]$, there is no information set that would involve nodes not present in this subgame — neither pursuer nor evader forget that s_p^1 was played. The optimal behavior in these subgames therefore depends only on the belief in $I_p[s_p^1]$, which is fixed due to the fixed behavior in the first round. We can therefore compute value of the subgame below $I_p[s_p^1]$ separately by making chance simulate the belief in this information set.

Let us construct a game $G[s_p^1]$ which consists of the information set $I_p[s_p^1]$ and the subgame beneath it. In this game, information set $I_p[s_p^1]$ is reached with probability $\beta = \sum_{s_e^1 \notin s_p^1} b_{\pi_e}(s_e^1)$, while with probability $1 - \beta$ the pursuer gets utility γ without play — this accounts for the reward the pursuer gets if he catches the evader in the first round by playing action s_p^1. The nature player simulates the belief $b[s_p^1]$ in the information set $I_p[s_p^1]$, so that the probability of every history in this information set, given this information set was reached, is identical with the original game. The value of the game $G[s_p^1]$ corresponds to the following part of the Equation (4):

$$\underbrace{\gamma \sum_{s_e^1 \in s_p^1} b_{\pi_e}(s_e^1)}_{\substack{\text{Evader caught} \\ \text{in the first round}}} + \underbrace{\left[\sum_{s_e^1 \notin s_p^1} b_{\pi_e}(s_e^1)\right]}_{\substack{\text{Evader not caught} \\ \text{in the first round}}} \sum_{s_e^1 \notin s_p^1} \sum_{s_e^0 \notin s_p^0} \underbrace{\left[\frac{b(s_e^0) \cdot \pi_e(s_e^0, s_e^1)}{\sum_{\tilde{s}_e^1 \notin s_p^1} \sum_{\tilde{s}_e^0 \notin s_p^0} b(\tilde{s}_e^0) \cdot \pi_e(\tilde{s}_e^0, \tilde{s}_e^1)} \cdot u(s_e^0 s_p^1 s_e^1)\right]}_{\text{Belief } b[s_p^1] \text{ of history } s_e^0 s_p^1 s_e^1 \text{ in } I_p[s_p^1]} \quad (5)$$

In the case of the game $G[s_p^1]$, there are multiple histories for every *current* position of the evader s_e^1 in the information set $I_p[s_p^1]$ (resulting from different initial locations of the evader s_e^0). We show that it is not necessary to account for different initial positions of the evader s_e^0, and thus all histories in $I_p[s_p^1]$ having the same current position of the evader s_e^1 can be merged. The resulting game contains a single history for each s_e^1 in $I_p[s_p^1]$, and thus this game is equivalent to a shorter horizon game $G^{t-1}\langle s_p^1, b_{\pi_e}\rangle$ up to multiplication of the utilities by γ to account for a round that has already passed. This allows using the solution of $G^{t-1}\langle s_p^1, b_{\pi_e}\rangle$ represented by value functions v^{t-1} to express the value of $G[s_p^1]$.

Definition 3. *Two deterministic game trees over nodes $\mathcal{H}_1, \mathcal{H}_2$ are isomorphic if there exists a bijection $\xi : \mathcal{H}_1 \to \mathcal{H}_2$ such that $v \in \mathcal{H}_1$ is a successor of $u \in \mathcal{H}_1$ if and only if $\xi(v)$ is a successor of $\xi(u)$, $n \in \mathcal{H}_1$ is a pursuer's node if and only if $\xi(n)$ is a pursuer's node, it is a terminal node if and only if $\xi(n)$ is a terminal node and the utilities $u(n) = u(\xi(n))$. Moreover the trees have the same informational structure so that two nodes $u, v \in \mathcal{H}_1$ are in the same information set if and only if nodes $\xi(u), \xi(v)$ are in the same information set.*

We can observe that subtrees of nodes $s_e^0 s_p^1 s_e^1$ and $\overline{s}_e^0 s_p^1 s_e^1$ (where s_e^0 and $\overline{s}_e^0$ stands for two different initial positions of the evader) are isomorphic as we can establish a bijection $\xi(s_e^0 s_p^1 s_e^1 h_{rest}) = \overline{s}_e^0 s_p^1 s_e^1 h_{rest}$. The utility of terminal histories does not depend on the initial position of the evader (only on the time when the pursuer managed to capture the evader). Whenever pursuer's node u is in information set I_p, node $\xi(u)$ is in I_p as well (because pursuer has no way to detect the evader's initial position). Moreover whenever evader cannot distinguish between two histories $s_e^0 s_p^1 s_e^1 \cdots s_p^q$ and $s_e^0 s_p^1 s_e^1 \cdots \overline{s}_p^q$, she cannot distinguish between histories $\overline{s}_e^0 s_p^1 s_e^1 \cdots s_p^q$ and $\overline{s}_e^0 s_p^1 s_e^1 \cdots \overline{s}_p^q$ either (because her uncertainty is related to the pursuer's move at round q, which does not depend on the initial position of the evader). Thus the subtrees have also the same informational structure.

Lemma 3. *Let I be the topmost information set of game $G[s_p^1]$ and let the belief $b[I]$ over nodes from I be known and fixed. Let $n_1, n_2 \in I$ be two nodes whose subtrees are isomorphic. Then a game G' with the same structure as G with any belief $b'[I]$ in I, satisfying $b[n_1] + b[n_2] = b'[n_1] + b'[n_2]$ and $b[n] = b'[n]$ for all nodes other than n_1 and n_2, has the same value as G.*

Thanks to the Lemma 3 and the isomorphism of the subtrees beneath $s_e^0 s_p^1 s_e^1$ and $\overline{s}_e^0 s_p^1 s_e^1$, histories $s_e^0 s_p^1 s_e^1$ and $\overline{s}_e^0 s_p^1 s_e^1$ can be merged and associated beliefs added up. By repeating this process, we end up with a single history for each current position of the evader s_e^1 (let $s_e^0 s_p^1 s_e^1$ be such history), whose belief is

$$b'[s_p^1](s_e^0 s_p^1 s_e^1) \coloneqq \frac{\sum_{s_e^0 \notin s_p^0} b(s_e^0) \cdot \pi_e(s_e^0, s_e^1)}{\sum_{\tilde{s}_e^1 \notin s_p^1} \sum_{\tilde{s}_e^0 \notin s_p^0} b(\tilde{s}_e^0) \cdot \pi_e(\tilde{s}_e^0, \tilde{s}_e^1)} = \frac{b_{\pi_e}(s_e^1)}{\sum_{\tilde{s}_e^1 \notin s_p^1} b_{\pi_e}(s_e^1)}; \quad b'[s_p^1](s_e^1) \text{ for short} \quad (6)$$

The updated belief $b'[s_p^1]$ from Equation (6) complies with belief b_{π_e} (computed according to Equation (3)) updated with the information that the evader is located in none of the vertices of s_p^1. The belief in $I_p[s_p^1]$ is identical with the belief in top-level information set of the game $G^{t-1}\langle s_p^1, b_{\pi_e}\rangle$; and hence the resulting game is identical to the game $G^{t-1}\langle s_p^1, b_{\pi_e}\rangle$ up to the multiplication by γ. The value of the game $G[s_p^1]$ (Equation (5)), from which this game was derived, is thus $\gamma v^{t-1}\langle s_p^1\rangle(b_{\pi_e})$. We substitute this value to Equation (4) to obtain

$$u(\emptyset) = \sum_{s_e^0 \in s_p^0} b(s_e^0) + \left[\sum_{s_e^0 \notin s_p^0} b(s_e^0)\right] \cdot \sum_{s_p^1} \pi_p(s_p^1) \cdot \left(\gamma v^t \langle s_p^1\rangle (b_{\pi_e})\right) \tag{7}$$

By allowing the players to choose their optimal one-step strategies π_p and π_e in Equation (7), we obtain the desired maximin formula shown in Equation (2). □

3.1 From value functions to optimal one-step strategies

The evaluation of the maximin formula from Equation (2) involves computation of optimal strategies of the players. In this section we show that if the value functions v^{t-1} are piecewise linear and convex functions represented by sets of α-vectors (which holds due to Theorem 2), the strategies can be found out by means of linear programming.

Due to the space constraints, we provide only the linear program for computing optimal one-step strategy of the pursuer in the game $G^t\langle s_p^0, b\rangle$ (its dual for computing evader's strategy can be found in the Appendix). At the beginning of each round, the pursuer realizes what vertices the evader is *not* located in, and hence updates his belief about the position of the evader. We can therefore restrict ourselves to the case where $b(s_e) = 0$ for all $s_e \in s_p^0$.

In the following linear program, the pursuer seeks for a strategy maximizing his expected utility against the best-responding opponent. He assumes strategies of the form "move to s_p^1 first and then follow strategy whose value is represented by $\alpha \in v^{t-1}\langle s_p^1\rangle$". The choice of α uniquely defines such strategy. The probability of playing each strategy $\alpha \in v^{t-1}\langle s_p^1\rangle$ is represented by variable $\hat{\pi}_p(s_p^1, \alpha)$. Constraint (9) corresponds to the value of playing such randomized strategy against the best-responding evader who starts in vertex s_e ($\alpha(s_e')$ denotes the value of α evaluated at pure belief corresponding to action s_e' of the evader). The evader starts in s_e with probability $b(s_e)$, hence the objective (8) calculates the expectation over individual $v(s_e)$. For the resulting one-step strategy of the pursuer, it holds that $\pi(s_p^1) = \sum_{\alpha \in v^{t-1}\langle s_p^1\rangle} \hat{\pi}(s_p^1, \alpha)$.

$$\max_{v, \hat{\pi}_p} \gamma \sum_{s_e \in \mathcal{V}} b(s_e) \cdot v(s_e) \tag{8}$$

$$\text{s.t.} \quad \sum_{s_p^1 \in \mathrm{adj}(s_p^0)\ ;\ \alpha \in v^{t-1}\langle s_p^1\rangle} \alpha(s_e') \cdot \hat{\pi}_p(s_p^1, \alpha) \geq v(s_e) \qquad \forall s_e \in \mathcal{V}\ \forall s_e' \in \mathrm{adj}(s_e) \tag{9}$$

$$\sum_{s_p^1 \in \mathrm{adj}(s_p^0)\ ;\ \alpha \in v^{t-1}\langle s_p^1\rangle} \hat{\pi}_p(s_p', \alpha) = 1 \tag{10}$$

$$\hat{\pi}_p(s_p^1, \alpha) \geq 0 \qquad \forall s_p^1 \in \mathrm{adj}(s_p^0)\ \forall \alpha \in v^{t-1}\langle s_p^1\rangle \tag{11}$$

3.2 Computing value functions

In each iteration of our value iteration algorithm, value functions v^t are constructed from the solution from the previous iteration — value functions v^{t-1}. By repeating this construction, a sequence of finite-horizon value functions $\{v^t\}_{t=0}^{\infty}$ approaching the values of the infinite-horizon game is being constructed. The value functions v^t that are about to be constructed, as well as v^t, are piecewise linear and convex (Theorem 2). In this section, we show that this allows us to avoid evaluating the dynamic programming operator H (Equation (2)) in every point in the belief space and enables us to construct v^t by considering only a finite subset of beliefs, corresponding to the extreme points of line segments of v^t. We proceed in two steps: (1) firstly we compute a function $Q^t_{\pi_p}\langle s_p^0\rangle$ corresponding to the expected utility the pursuer gets if he plays π_p at the first round of the longer horizon game $G^t\langle s_p^0, b\rangle$; (2) then we show how to compute $v^t\langle s_p^0\rangle$ as a combination of multiple $Q^t_{\pi_p}\langle s_p^0\rangle$ for properly chosen one-step strategies π_p. We start with a formal definition of function $Q^t_{\pi_p}\langle s_p^0\rangle$.

Definition 4. *Let π_p be pursuer's one-step strategy for the first round of the game $G^t\langle s_p^0, b\rangle$. The* value of π_p *is a function $Q^t_{\pi_p}\langle s_p^0\rangle$ assigning the expected reward the pursuer gets in the game $G^t\langle s_p^0, b\rangle$ against the best-responding opponent, when he plays π_p in the first round and continues by playing according to his optimal strategy in the rest of the game, i.e.*

$$Q^t_{\pi_p}\langle s_p^0\rangle(b) \coloneqq \sum_{s_e \in s_p^0} b(s_e) + \gamma \left[\sum_{s_e \in \mathcal{V}\setminus s_p^0} b(s_e)\right] \cdot \min_{\pi_e} \sum_{s_p^1 \in \mathcal{V}^N} \pi_p(s_p^1) \cdot v^{t-1}\langle s_p^1\rangle(b_{\pi_e}) \tag{12}$$

According to the previous definition, once the first round of the game is over, the pursuer continues with his optimal strategy. The following lemma shows that this optimal strategy for the rest of the game is well-defined and its value is characterized by the value functions v^{t-1}.

Lemma 4. *Let π_p be pursuer's fixed one-step strategy for the first round of the game. For every belief b there are strategies $\sigma_p[s_p^1]$, one for each $s_p^1 \in \mathrm{adj}(s_p^0)$, represented by α-vectors $\alpha[s_p^1] \in v^{t-1}\langle s_p^1\rangle$, such that it is optimal to follow $\sigma_p[s_p^1]$ when s_p^1 was played in the first round of the game. The value of strategy σ_p prescribing the pursuer to play according to π_p in the first round and continue by using respective $\sigma_p[s_p^1]$ is linear and the corresponding α-vector satisfies*

$$\alpha^{\sigma_p}(s_e) = \begin{cases} 1 & s_e \in s_p \\ \gamma \min\limits_{s'_e \in \mathrm{adj}(s_e)} \sum\limits_{s_p^1} \pi_p(s_p^1) \cdot \alpha[s_p^1](s'_e) & \text{otherwise} \end{cases} \tag{13}$$

Lemma 4 gives us a direct algorithm for computing $Q^t_{\pi_p}$. PWLC functions v^{t-1} correspond to a finite number of horizon-t strategies, represented by a finite number of α-vectors. Thus there is only a finite number of ways to choose

$\hat{v}^t\langle s_p^0\rangle \leftarrow \left\{\mathbf{0}^{|\mathcal{V}|}\right\}$
$\hat{\Pi}_p = \emptyset$
while $\exists b \in \Delta(\mathcal{V}), \pi_p \notin \Pi_p : Q^t_{\pi_p}\langle s_p^0\rangle(b) > \hat{v}^t\langle s_p^0\rangle(b)$ **do**
  $\pi_p \leftarrow$ optimal strategy of the pursuer at belief b for the first round (see (8))
  $\hat{\Pi}_p \leftarrow \hat{\Pi}_p \cup \{\pi_p\}$
  $\hat{v}^t\langle s_p^0\rangle \leftarrow \hat{v}^t\langle s_p^0\rangle \oplus Q^t_{\pi_p}\langle s_p\rangle$
return $\hat{v}^t\langle s_p\rangle$

Algorithm 1: Incremental construction of value function $v^t\langle s_p\rangle$

strategies $\sigma_p[s_p^1]$ from Lemma 4, which can be found by means of enumeration. The maximization over linear functions representing value of such strategies corresponds to the function $Q^t_{\pi_p}\langle s_p\rangle$ which is thus piecewise linear and convex.

The definition of $Q^t_{\pi_p}\langle s_p\rangle$ implies that we can compute the value function $v^{t+1}\langle s_p\rangle$ by allowing the pursuer to play arbitrary strategy π_p, when

$$v^t\langle s_p^0\rangle(b) = \max_{\pi_p} Q^t_{\pi_p}\langle s_p^0\rangle(b) \tag{14}$$

As a consequence of Theorem 1, it is sufficient to consider a finite set Π_p of strategies in the maximizer of Equation (14) and thus obtain $v^t\langle s_p^0\rangle$ as the pointwise maximum from respective $Q^t_{\pi_p}\langle s_p^0\rangle$ functions, $v^t\langle s_p\rangle = \bigoplus_{\pi_p\in\Pi_p} Q^t_{\pi_p}\langle s_p\rangle$. The set of such strategies Π_p is however initially unknown. We propose an algorithm (Algorithm 1) that constructs both the set of strategies $\hat{\Pi}_p$ and the value function $\hat{v}^t\langle s_p^0\rangle$ incrementally by iteratively verifying whether the current set of the strategies $\hat{\Pi}_p$ is sufficient for obtaining the actual value function $v^t\langle s_p^0\rangle$.

The Algorithm 1 is constructing a set of strategies $\hat{\Pi}_p$ and a corresponding estimate of value function $\hat{v}^t\langle s_p^0\rangle = \bigoplus_{\pi_p\in\hat{\Pi}_p} Q^t_{\pi_p}\langle s_p^0\rangle$, starting with an empty set $\hat{\Pi}_p$. At each iteration, it verifies whether strategies $\hat{\Pi}_p$ used to compute current $\hat{v}^{t+1}\langle s_p^0\rangle$ are optimal in every belief $b \in \Delta(\mathcal{V})$. If it finds a belief b where the strategy can be improved, i.e. there exists π_p such that $Q^t_{\pi_p}\langle s_p^0\rangle(b) > \hat{v}^t\langle s_p^0\rangle(b)$, it updates the set $\hat{\Pi}_p$ and recomputes $\hat{v}^t\langle s_p\rangle$. If no such belief is found, all required strategies were considered and $\hat{v}^t\langle s_p^0\rangle = v^t\langle s_p^0\rangle$.

Whenever the value function $\hat{v}^t\langle s_p^0\rangle$ is not yet optimal in the whole belief space, i.e. there exists a belief b where $Q^t_{\pi_p}\langle s_p^0\rangle(b) > \hat{v}^t\langle s_p^0\rangle(b)$, there exists a belief b' with the same property that forms an extreme point of a line segment on $\hat{v}^t\langle s_p^0\rangle$. This property is characterized by Lemma 5.

Lemma 5. *If there is a belief b where $v^t\langle s_p^0\rangle(b) > \hat{v}^t\langle s_p^0\rangle(b)$, there must be a belief b' that forms an extreme point of a line segment on the surface of $\hat{v}^t\langle s_p^0\rangle$ where $v^t\langle s_p^0\rangle(b') > \hat{v}^t\langle s_p^0\rangle(b')$.*

Thanks to Lemma 5, we can consider only a *finite* set of beliefs that form extreme points of line segments on the value function $\hat{v}^t\langle s_p^0\rangle$. In every iteration, a one-step strategy that is optimal at some belief point (and thus must be present in Π_p) is added to the set $\hat{\Pi}_p$. As a consequence of the Theorem 1, the set Π_p

that is necessary to obtain the optimal value function $v^t\langle s_p^0\rangle$ is finite. Hence after a finite number of iterations, the Algorithm 1 terminates.

3.3 Convergence and uniqueness of the solution

We demonstrate the convergence properties of our value iteration algorithm and the uniqueness of the value functions solving the infinite horizon game by showing that the dynamic programming operator H (Equation 2) is a contraction mapping. The desired properties then follows from the Banach's fixed point theorem [3]. We show the contractivity of H under the following max-norm:

$$\|v - \overline{v}\| = \max_{s_p^0 \in \mathcal{V}^N} \max_{b \in \Delta(\mathcal{V})} |v\left\langle s_p^0 \right\rangle(b) - \overline{v}\left\langle s_p^0 \right\rangle(b)| \tag{15}$$

Lemma 6. *The operator H is a contraction with contractivity factor $\gamma < 1$ under max-norm.*

Theorem 4. *There is a unique set of value functions v^* satisfying $v^* = Hv^*$ and the recursive application of H converges to v^*. Series $\{v^t\}_{i=0}^{\infty}$ thus converges to value functions of an infinite horizon game.*

Proof. The operator H is a contraction mapping defined on a metric space of sets of bounded functions defined on the belief space. By applying Banach's fixed point theorem [3] we get that H has a unique fixed point v^* and the recursive application of H converges to v^*. □

Proposition 1. *After t iterations of the value iteration algorithm, the value function v^t is γ^t-optimal (i.e. $\|v^t - v^*\| \leq \gamma^t$).*

4 Conclusion

We present the first algorithm for solving the class of two-player discounted pursuit-evasion games with infinite horizon and partial observability, where the evader is assumed to be perfectly informed about the current state of the game (i.e. position of pursuer's units). This class of games has a significant relevance in security domains where a robust strategy that provides guarantees in the worst case is often desirable.

Our algorithm is a modification of the well-known value iteration algorithm for solving Partially Observable Markov Decision Processes (POMDPs), or stochastic games with concurrent moves. We show that the strategies can be compactly represented using value functions that depend on the location of the pursuing units and the belief about the position of the evader, but not explicitly on the history of moves. These value functions are piecewise linear and convex and allow us to design a dynamic programming operator for the value iteration algorithm.

Our work is the first step towards many practical algorithms for solving discounted stochastic games with one-sided partial observability. These can be applied in many scenarios requiring robust strategies and thus our work opens the whole new area of research in algorithmic and computational game theory. One natural continuation is an adaptation of point-based approximation algorithms for POMDPs to improve the scalability of the value iteration algorithm.

References

1. Chatterjee, K., Doyen, L.: Partial-observation stochastic games: How to win when belief fails. ACM Transactions on Computational Logic (TOCL) 15(2), 16 (2014)
2. Chung, T.H., Hollinger, G.A., Isler, V.: Search and pursuit-evasion in mobile robotics. Autonomous robots 31(4), 299–316 (2011)
3. Ciesielski, K., et al.: On Stefan Banach and some of his results. Banach Journal of Mathematical Analysis 1(1), 1–10 (2007)
4. Hansen, E.A., Bernstein, D.S., Zilberstein, S.: Dynamic programming for partially observable stochastic games. In: AAAI. vol. 4, pp. 709–715 (2004)
5. Koller, D., Megiddo, N., Von Stengel, B.: Efficient computation of equilibria for extensive two-person games. Games and Economic Behavior 14(2), 247–259 (1996)
6. McEneaney, W.M.: Some classes of imperfect information finite state-space stochastic games with finite-dimensional solutions. Applied Mathematics and Optimization 50(2), 87–118 (2004)
7. Monahan, G.E.: State of the arta survey of partially observable Markov decision processes: theory, models, and algorithms. Management Science 28(1), 1–16 (1982)
8. Pineau, J., Gordon, G., Thrun, S., et al.: Point-based value iteration: An anytime algorithm for POMDPs. In: IJCAI. vol. 3, pp. 1025–1032 (2003)
9. Shapley, L.S.: Stochastic games. Proceedings of the National Academy of Sciences 39(10), 1095–1100 (1953)
10. Smallwood, R.D., Sondik, E.J.: The optimal control of partially observable Markov processes over a finite horizon. Operations Research 21(5), 1071–1088 (1973)
11. Smith, T., Simmons, R.: Point-based POMDP algorithms: Improved analysis and implementation. arXiv preprint arXiv:1207.1412 (2012)
12. Vanderbei, R.J.: Linear programming. Springer (2014)
13. Vidal, R., Shakernia, O., Kim, H.J., Shim, D.H., Sastry, S.: Probabilistic pursuit-evasion games: theory, implementation, and experimental evaluation. Robotics and Automation, IEEE Transactions on 18(5), 662–669 (2002)

APPENDIX

A Proofs

A.1 Proof of Lemma 2

The expected reward in the root node $u(\emptyset)$ is

$$u(\emptyset) = \overbrace{\sum_{s_e^0 \in s_p^0} b(s_e^0)}^{\text{Evader is caught}} + \overbrace{\gamma \sum_{s_e^0 \notin s_p^0} \sum_{s_p^1} \sum_{s_e^1 \in s_p^1} b(s_e^0) \cdot \pi_p(s_p^1) \cdot \pi_e(s_e^0, s_e^1)}^{\text{Evader is caught in the first round}} + \underbrace{\qquad\qquad}_{\text{All terminal histories for the first round, i.e. shorter than two actions}}$$

$$+ \underbrace{\sum_{s_e^0 \notin s_p^0} \sum_{s_p^1} \sum_{s_e^1 \notin s_p^1} b(s_e^0) \cdot \pi_p(s_p^1) \cdot \pi_e(s_e^0, s_e^1) \cdot u(s_e^0 s_p^1 s_e^1)}_{\text{Evader is not caught in the first round}}$$

The derivation of Equation (4) is then just a technical derivation involving operations with sums and normalization of conditional probability distributions. Throughout the derivation we will use the equation of the transformed belief (Equation (3)).

$$\begin{aligned} u(\emptyset) &= \sum_{s_e^0 \in s_p^0} b(s_e^0) + \gamma \sum_{s_e^0 \notin s_p^0} \sum_{s_p^1} \sum_{s_e^1 \in s_p^1} b(s_e^0) \cdot \pi_p(s_p^1) \cdot \pi_e(s_e^0, s_e^1) \; + \\ &\quad + \sum_{s_e^0 \notin s_p^0} \sum_{s_p^1} \sum_{s_e^1 \notin s_p^1} b(s_e^0) \cdot \pi_p(s_p^1) \cdot \pi_e(s_e^0, s_e^1) \cdot u(s_e^0 s_p^1 s_e^1) \\ &= \sum_{s_e^0 \in s_p^0} b(s_e^0) + \gamma \sum_{s_p^1} \pi_p(s_p^1) \sum_{s_e^1 \in s_p^1} \sum_{s_e^0 \notin s_p^0} b(s_e^0) \cdot \pi_e(s_e^0, s_e^1) \; + \\ &\quad + \sum_{s_p^1} \pi_p(s_p^1) \sum_{s_e^1 \notin s_p^1} \sum_{s_e^0 \notin s_p^0} b(s_e^0) \cdot \pi_e(s_e^0, s_e^1) \cdot u(s_e^0 s_p^1 s_e^1) \end{aligned}$$

$$= \sum_{s_e^0 \in s_p^0} b(s_e^0) + \gamma \left[\sum_{s_e^0 \notin s_p^0} b(s_e^0) \right] \sum_{s_p^1} \pi_p(s_p^1) \sum_{s_e^1 \in s_p^1} \frac{\sum_{s_e^0 \notin s_p^0} b(s_e^0) \cdot \pi_e(s_e^0, s_e^1)}{\sum_{s_e^0 \notin s_p^0} b(s_e^0)} \; +$$

$$+ \sum_{s_p^1} \pi_p(s_p^1) \cdot \left[\sum_{\tilde{s}_e^1 \notin s_p^1} \sum_{\tilde{s}_e^0 \notin s_p^0} b(s_e^0) \cdot \pi_e(s_e^0, s_e^1) \right] \sum_{s_e^1 \notin s_p^1} \sum_{s_e^0 \notin s_p^0} \left[\frac{b(s_e^0) \cdot \pi_e(s_e^0, s_e^1)}{\sum_{\tilde{s}_e^1 \notin s_p^1} \sum_{\tilde{s}_e^0 \notin s_p^0} b(s_e^0) \cdot \pi_e(s_e^0, s_e^1)} \cdot u(s_e^0 s_p^1 s_e^1) \right]$$

$$= \sum_{s_e^0 \in s_p^0} b(s_e^0) + \left[\sum_{s_e^0 \notin s_p^0} b(s_e^0) \right] \left(\gamma \sum_{s_p^1} \pi_p(s_p^1) \sum_{s_e^1 \in s_p^1} \frac{\sum\limits_{s_e^0 \notin s_p^0} b(s_e^0) \cdot \pi_e(s_e^0, s_e^1)}{\sum\limits_{s_e^0 \notin s_p^0} b(s_e^0)} \; + \right.$$

$$\left. + \sum_{s_p^1} \pi_p(s_p^1) \left[\sum_{\tilde{s}_e^1 \notin s_p^1} \frac{\sum\limits_{\tilde{s}_e^0 \notin s_p^0} b(s_e^0) \cdot \pi_e(s_e^0, s_e^1)}{\sum\limits_{s_e^0 \notin s_p^0} b(s_e^0)} \right] \sum_{s_e^1 \notin s_p^1} \sum_{s_e^0 \notin s_p^0} \left[\frac{b(s_e^0) \cdot \pi_e(s_e^0, s_e^1)}{\sum\limits_{\tilde{s}_e^1 \notin s_p^1} \sum\limits_{\tilde{s}_e^0 \notin s_p^0} b(s_e^0) \cdot \pi_e(s_e^0, s_e^1)} \cdot u(s_e^0 s_p^1 s_e^1) \right] \right)$$

$$= \sum_{s_e^0 \in s_p^0} b(s_e^0) + \left[\sum_{s_e^0 \notin s_p^0} b(s_e^0) \right] \left(\gamma \sum_{s_p^1} \pi_p(s_p^1) \sum_{s_e^1 \in s_p^1} b_{\pi_e}(s_e^1) \; + \right.$$

$$\left. + \sum_{s_p^1} \pi_p(s_p^1) \left[\sum_{s_e^1 \notin s_p^1} b_{\pi_e}(s_e^1) \right] \sum_{s_e^1 \notin s_p^1} \sum_{s_e^0 \notin s_p^0} \left[\frac{b(s_e^0) \cdot \pi_e(s_e^0, s_e^1)}{\sum\limits_{\tilde{s}_e^1 \notin s_p^1} \sum\limits_{\tilde{s}_e^0 \notin s_p^0} b(\tilde{s}_e^0) \cdot \pi_e(\tilde{s}_e^0, \tilde{s}_e^1)} \cdot u(s_e^0 s_p^1 s_e^1) \right] \right)$$

$$= \sum_{s_e^0 \in s_p^0} b(s_e^0) + \left[\sum_{s_e^0 \notin s_p^0} b(s_e^0) \right] \cdot \sum_{s_p^1} \pi_p(s_p^1) \left(\gamma \sum_{s_e^1 \in s_p^1} b_{\pi_e}(s_e^1) \; + \right.$$

$$\left. + \left[\sum_{s_e^1 \notin s_p^1} b_{\pi_e}(s_e^1) \right] \sum_{s_e^1 \notin s_p^1} \sum_{s_e^0 \notin s_p^0} \left[\frac{b(s_e^0) \cdot \pi_e(s_e^0, s_e^1)}{\sum\limits_{\tilde{s}_e^1 \notin s_p^1} \sum\limits_{\tilde{s}_e^0 \notin s_p^0} b(\tilde{s}_e^0) \cdot \pi_e(\tilde{s}_e^0, \tilde{s}_e^1)} \cdot u(s_e^0 s_p^1 s_e^1) \right] \right)$$

A.2 Proof of Lemma 3

Proof. Let $\mathcal{H}_1$ be histories in the subtree of n_1, $\mathcal{H}_2$ be histories in the subtree of n_2 and $\xi : \mathcal{H}_1 \to \mathcal{H}_2$ be a bijection from the definition of the game tree isomorphism. Let σ be a Nash equilibrium strategy profile in the game G and $\sigma[n_1]$, $\sigma[n_2]$ be the optimal behaviors in subtrees of n_1 and n_2 when the strategy profile σ is followed.

First of all, we show that we can modify the behavior in the subtree of n_2 to $\sigma'[n_2]$, where $\sigma'[n_2](n) = \sigma[n_1](\xi(n))$ without changing the expected utility (and hence we can assume that the strategies in both subtrees are the same). Note that strategy profile $\sigma'[n_2]$ is valid in the subtree of n_2 due to the fact that

subtrees beneath n_1 and n_2 are isomorphic. Moreover this change does not affect behavior in the rest of the game, as $\sigma[n_1]$ was consistent with this behavior.

Assume that the strategy profile $\sigma'[n_2]$ is *not* optimal in the subtree of n_2. Strategies $\sigma[n_1]$ and $\sigma'[n_2]$ induce the same distribution over leaf nodes in subtrees of n_1 and n_2 (up to the bijection ξ) and these leaf nodes have the same utility values. Hence if one of the players wanted to deviate from $\sigma'[n_2]$ in the subtree of n_2, he would have wanted to do the same in the case of $\sigma[n_1]$ used in the subtree of n_1. Thus the strategy $\sigma'[n_2]$ must be an optimal behavior in the subtree of n_2.

Strategies $\sigma[n_1]$ and $\sigma'[n_2]$ induce the same distribution over the leafs in the respective subtrees (given that the node n_1, resp. n_2, was reached). We construct a game G' from G by modifying the probability of reaching nodes n_1 and n_2 (the belief in information set I of G' is denoted b'). The probability of reaching one of the terminal histories $u \in \mathcal{H}_1$ and $\xi(u)$ in G' remains the same as in G as long as $b'[I](n_1) + b'[I](n_2) = b[I](n_1) + b[I](n_2)$. The modification of the belief thus does not change the optimal behavior of the players, and thus does not change the value of the game.

A.3 Proof of Lemma 4

Assume that the pursuer played action s_p^1 in the first round of the game (drawn from π_p). The game moves to a shorter horizon game $G^{t-1}\langle s_p^1, b_{\pi_e}\rangle$ with the belief updated according to Equation (3) (where π_e is evader's best response to π_p). The optimal strategy $\sigma_p[s_p^1]$ of the pursuer in $G^{t-1}\langle s_p^1, b_{\pi_e}\rangle$ is optimal at belief b_{π_e}, hence, by definition, its corresponding α-vector $\alpha[s_p^1]$ is present in the value function $v^{t-1}\langle s_p^1\rangle$ which expresses the value of all non-dominated strategies in the game $G^{t-1}\langle s_p^1, b_{\pi_e}\rangle$.

The value of every strategy is linear in belief (Lemma 1). It is therefore sufficient to define the value of the strategy σ_p in each of the pure beliefs to form the α-vector α^{σ_p}. If the evader is located in the same vertex as one of the units of the pursuer, the game ends immediately and the pursuer gets utility $\gamma^0 = 1$. If the evader is not immediately caught, she chooses a vertex s'_e adjacent to her current position s_e so that the expected utility of the pursuer is minimized. We know that if the pursuer plays s_p^1 in the first round and follows with $\sigma_p[s_p^1]$ afterwards, the expected utility of the pursuer is represented by the α-vector $\alpha_p[s_p^1]$ evaluated at the pure belief corresponding to the new position of the evader s'_e, multiplied by γ as one round has already passed.

A.4 Proof of Lemma 5

We show that if there is a belief point b, where the value function $\hat{v}^t\langle s_p^0\rangle$ can be improved, the value function $v^t\langle s_p^0\rangle$ can also be improved in at least one of the vertices of the facet on the surface of $v^t\langle s_p^0\rangle$. We prove this by contradiction. Let F be the facet on the surface of $\hat{v}^t\langle s_p^0\rangle$ that the belief b is projected on (i.e. b is a

convex combination of coordinates of vertices of F). Assume that the value function $\hat{v}^t\langle s_p^0\rangle$ can be improved in b, but not in *any* of the vertices of F. This means that for every belief b' corresponding to the extreme point of the facet F, the value function $\hat{v}^t\langle s_p^0\rangle$ is optimal, i.e. $\hat{v}^t\langle s_p^0\rangle(b') = v^t\langle s_p^0\rangle(b')$. Moreover, as the value function can be improved at belief b, it holds that $\hat{v}^t\langle s_p^0\rangle(b) < v^t\langle s_p^0\rangle(b)$. This means that the value of $v^t\langle s_p^0\rangle$ at b is above facet F, as the original value $\hat{v}^t\langle s_p^0\rangle(b)$ was a convex combination of its vertices, which contradicts the convexity of the value function (Theorem 2).

A.5 Proof of Lemma 6

Let us define $Q^v_{\pi_p}\langle s_p\rangle$ similarly to $Q^t_{\pi_p}\langle s_p\rangle$ used in Section 3.2, with the exception that we refer to an arbitrary value functions v instead of v^{t-1}:

$$Q^v_{\pi_p}\langle s_p\rangle(b) \coloneqq \sum_{s_e\in s_p} b(s_e) + \gamma\left[\sum_{s_e\in\mathcal{V}\setminus s_p} b(s_e)\right]\cdot \min_{\pi_e}\sum_{s'_p\in\mathcal{V}^N}\pi_p(s'_p)\cdot v\left\langle s'_p\right\rangle(b_{\pi_e}) \tag{16}$$

The proof will closely follow the structure of the proof of Theorem 1 in [11]. First of all we show that for every $s_p\in\mathcal{V}^N$ and every valid pursuer's one-step strategy π_p, the mapping $v\mapsto Q^v_{\pi_p}\langle s_p\rangle$ has a contractivity factor γ, then by inspecting all possible s_p and π_p we show that the same holds for H. Note that the difference $|Q^v_{\pi_p}\langle s_p\rangle(b) - Q^{\overline{v}}_{\pi_p}\langle s_p\rangle(b)|$ is maximized if b is chosen so that the evader is not initially caught according to b (i.e. $\sum_{s_e\in s_p} b(s_e)=0$), which allows us to simplify the derivation.

$$\begin{aligned}
&\left\|Q^v_{\pi_p}\langle s_p\rangle - Q^{\overline{v}}_{\pi_p}\langle s_p\rangle\right\| = \max_b\left|Q^v_{\pi_p}\langle s_p\rangle(b) - Q^{\overline{v}}_{\pi_p}\langle s_p\rangle(b)\right|\\
&= \max_b\left|\gamma\min_{\pi_e}\textstyle\sum_{s'_p}\pi_p(s'_p)\cdot v\left\langle s'_p\right\rangle(b_{\pi_e}) - \gamma\min_{\pi'_e}\sum_{s'_p}\pi_p(s'_p)\cdot\overline{v}\left\langle s'_p\right\rangle(b_{\pi'_e})\right|\\
&= \gamma\max_b\left|\min_{\pi_e}\textstyle\sum_{s'_p}\pi_p(s'_p)\cdot v\left\langle s'_p\right\rangle(b_{\pi_e}) - \min_{\pi'_e}\sum_{s'_p}\pi_p(s'_p)\cdot\overline{v}\left\langle s'_p\right\rangle(b_{\pi'_e})\right|\\
&\le \gamma\max_b\max_{\pi_e}\left|\textstyle\sum_{s'_p}\pi_p(s'_p)\cdot v\left\langle s'_p\right\rangle(b_{\pi_e}) - \sum_{s'_p}\pi_p(s'_p)\cdot\overline{v}\left\langle s'_p\right\rangle(b_{\pi_e})\right|\\
&\le \gamma\max_{b'}\left|\textstyle\sum_{s'_p}\pi_p(s'_p)\cdot\left[v\left\langle s'_p\right\rangle(b') - \overline{v}\left\langle s'_p\right\rangle(b')\right]\right|\\
&\le \gamma\max_{b'}\textstyle\sum_{s'_p}\pi_p(s'_p)\cdot\left|v\left\langle s'_p\right\rangle(b') - \overline{v}\left\langle s'_p\right\rangle(b')\right|\\
&\le \gamma\max_{b'}\textstyle\sum_{s'_p}\pi_p(s'_p)\cdot\|v-\overline{v}\|\\
&\le \gamma\|v-\overline{v}\|
\end{aligned}$$

Let us now choose arbitrary $s_p\in\mathcal{V}^N$ and $b\in\Delta(\mathcal{V})$. Without loss of generality, let us assume that $(Hv)\langle s_p\rangle(b) \ge (H\overline{v})\langle s_p\rangle(b)$. Let π_p^* be an optimal one-shot strategy in b w.r.t. v (i.e. maximizing $Q^v_{\pi_p^*}\langle s_p\rangle(b)$) and $\overline{\pi}_p$ an optimal strategy w.r.t. $\overline{v}$. Note that $(Hv)\langle s_p\rangle(b) = Q^v_{\pi_p^*}\langle s_p\rangle(b)$ and $(H\overline{v})\langle s_p\rangle(b) =$

$Q^{\overline{v}}_{\overline{\pi}_p}\langle s_p\rangle(b)$. It holds that $Q^{\overline{v}}_{\pi^*_p}\langle s_p\rangle(b) \leq Q^{\overline{v}}_{\overline{\pi}_p}\langle s_p\rangle(b) \leq Q^{v}_{\pi^*_p}\langle s_p\rangle(b)$. Then:

$$\begin{aligned}
|(Hv)\langle s_p\rangle(b) - (H\overline{v})\langle s_p\rangle(b)| &= \left|Q^{v}_{\pi^*_p}\langle s_p\rangle(b) - Q^{\overline{v}}_{\overline{\pi}_p}\langle s_p\rangle(b)\right| \\
&\leq \left|Q^{v}_{\pi^*_p}\langle s_p\rangle(b) - Q^{\overline{v}}_{\pi^*_p}\langle s_p\rangle(b)\right| \\
&\leq \max_{\pi_p}\left|Q^{v}_{\pi_p}\langle s_p\rangle(b) - Q^{\overline{v}}_{\pi_p}\langle s_p\rangle(b)\right| \\
&\leq \gamma \cdot \|v - \overline{v}\|
\end{aligned}$$

A.6 Proof of Proposition 1

The minimum reward the pursuer can get is zero, hence $v^0\langle s^0_p\rangle(b) \geq 0$ for every initial position of the pursuer s^0_p and every his belief b about the position of the evader. Similarly the maximum reward is one, hence $v^*\langle s^0_p\rangle(b) \leq 1$ for every s^0_p and every b. It holds therefore that $\|v^0 - v^*\| \leq 1$. Due to the contractivity factor γ of the dynamic operator H, after t iterations the distance $\|v^t - v^*\| \leq \gamma^t\|v^0 - v^*\| \leq \gamma^t$ which completes the proof.

B Computing optimal strategies

B.1 Computing evader's strategy

The idea behind the linear program for computing optimal evader's one-step strategy in game $G^t\langle s^0_p, b\rangle$ is similar to the one for solving the game from the perspective of the pursuer (Section 3.1). In this case the roles of the players are reversed — the evader seeks for her strategy, while the pursuer best-responds it. Similarly as in Section 3.1, we focus on the case where $b(s^0_e) = 0$ for all $s^0_e \in s^0_p$.

By choosing the strategy π_e, the evader decides the belief b_{π_e} in the shorter horizon game $G^{t-1}\langle s^1_p, b_{\pi_e}\rangle$ (see Equation (3)). The calculation of this updated belief is done by means of the Constraint (19). If the pursuer decides to move to the set of vertices s^1_p, his expected utility is described by the value function $v^{t-1}\langle s^1_p\rangle$ multiplied by γ to account for the first round of the game. The pursuer best-responds by choosing the best s^1_p from his perspective, i.e. the one that maximizes his utility, which is characterized by a set of best-response constraints (Equation (18)). Note that the maximization over α-vectors in the value functions is rewritten using a set of inequality constraints, one for each α-vector. The evader then seeks for a strategy that minimizes the expected utility of the pursuer (Equation (17)).

$$\min_{V,\pi_e,b_{\pi_e}} \quad V \tag{17}$$

$$\text{s.t. } \gamma \sum_{s'_e \in \mathcal{V}} \alpha(s'_e) \cdot b_{\pi_e}(s'_e) \leq V \qquad \forall s_p^1 \in (s_p^0) \;\; \forall \alpha \in v^{t-1}\langle s_p^1 \rangle \tag{18}$$

$$\sum_{s_e \in \mathcal{V} \setminus s_p^0} b(s_e) \cdot \pi_e(s_e, s'_e) = b_{\pi_e}(s'_e) \qquad \forall s'_e \in \mathcal{V} \tag{19}$$

$$\sum_{s'_e \in \text{adj}(s_e)} \pi_e(s_e, s'_e) = 1 \qquad \forall s_e \in \mathcal{V} \tag{20}$$

$$\pi_e(s_e, s'_e) \geq 0 \qquad \forall s_e \in \mathcal{V} \;\; \forall s'_e \in \text{adj}(s_e) \tag{21}$$